

IPECAC: AN EXPERT SYSTEM FOR THE MANAGEMENT OF POISONING INCIDENTS

IPECAC (Interactive Poison Expert for Classification And Control) is an expert software system for the diagnosis and management of poisoning that is being developed jointly by APL and the Maryland Poison Center. The goal is a microcomputer-based expert system that is inexpensive enough to be installed in hospital emergency rooms or doctors' offices and capable of handling 75 percent of all inquiries. Expert knowledge has been distilled from written material, from interviews with clinical experts, and from observing Maryland Poison Center staff members handling telephone inquiries. An innovative model for representing knowledge has resulted, and a pilot expert system has been written whose architecture implements the model. Clinical testing of the pilot system is now underway.

BACKGROUND

There are approximately 30 certified¹ regional poison centers in the United States. Each center answers calls from the general public, hospital emergency rooms, and private physicians in its area 24 hours a day. More centers are needed, but they are expensive to maintain because of the need for trained clinical staff members.

The Maryland Poison Center, located at the University of Maryland School of Pharmacy in Baltimore, is typical. It is directed by a doctor of pharmacy, and it employs ten people full time, a half-dozen part time, and five health-professional students. Two of the staff members have doctorate degrees, and the rest are graduate pharmacists and nurses, mostly at the bachelor's level. On the average, the center receives about 50,000 telephone inquiries per year; of these, there are about 30,000 cases of human exposure to potentially toxic substances.

With calls from the general public, the center tries to obtain as much information as possible and attempts to identify the substance(s) involved, the dosages, and symptoms. It then recommends a course of action that can range from doing nothing, up through observing, administering home treatment (e.g., ipecac syrup), or taking the patient to a hospital emergency room for more elaborate treatment. Inquiries from physicians usually seek specific information such as expected toxic effects or management strategies.

In about 95 percent of the cases, the substance is easily identified and the management strategy is clear. The remaining cases are handled by the clinical and pharmaceutical experience and judgment of the highly trained staff, who often consult with one another on unusual problems. The staff members report that they have acquired an important store of expert experience over the years.

The primary information source used by the center is a commercial product called Poisindex[®],² a database collected and maintained by the Micromedex Company of Denver. Micromedex also markets two other (and larger) databases called Drugdex[®] (for medications) and

Emergindex[®] (for hospital emergency rooms). Poisindex contains about 40 million bytes of information that includes 450,000 separate listings describing the toxic substances in about 92,000 different commercial products and 265 management procedures for different poisoning problems.

Poisindex is distributed to about 1100 subscribers on both microfiche and computer-compatible compact disk (see Fig. 1). The information is indexed in three ways: by the name of the commercial product, by physical description in the case of medicines, and by toxic substance for the 265 management strategies.

Poisindex is an invaluable source of information, but there are limitations on how it can be searched. For example, if a caller is not sure what type of pill a child has ingested but can describe symptoms the child is showing, Poisindex cannot be searched for a list of possible medications. In such cases, the clinical experience of the staff must be used.

THE GOAL

We do not seek to replace skilled clinicians, but we do hope to use their skills more efficiently. By develop-



Figure 1—The Poison Center staff relies heavily on the Poisindex Information System, which is stored on a compact disk and accessed from a personal computer.

ing an expert system that could handle the most routine inquiries, for example, the most common 75 percent, we would free the clinical staff to concentrate on the more complex problems from a larger population base.

A microcomputer-based system would be inexpensive enough to be installed in most hospital emergency rooms and perhaps in many doctors' offices. The expert system would know its limitations. When unusual substances were involved or when cases were extremely complex, the system would recommend that a regional poison center be called. Indeed, triage would be an important role for such a system.

CURRENT EFFORT

The current effort has been devoted to constructing a pilot system. We have chosen a clinically useful but circumscribed problem domain: ingestion of antihistamines or decongestants (see the boxed insert). In 1985, 5.5 percent of all toxic exposures involved these products.³

Expert knowledge has been gathered from published material such as *Poisindex*,² Maryland Poison Center management protocols, and standard textbooks.^{4,5,6} The clinical staff has been interviewed both informally and through structured questionnaires. In addition, the staff has been monitored as it handled inquiries. Study of these interactions has led to the scheme of knowledge representation described below.

THE CLINICAL STAFF AT WORK

Among the interactions monitored, the following three are both representative and instructive.

Case 1

A paramedic called. He had been summoned to the scene of a poisoning and was calling from the house. All he knew was that the subject had drunk something from a small bottle that had no label, but had the letters SKF on the lid. The residue in the bottle was a clear, sweet-smelling liquid. The call was taken by one of the experts who placed the call on hold and polled two or three pharmacists who were present. One immediately said "Compazine," a comment that was followed by one or more suggestions from the others. The list of possible drugs was transmitted to the paramedic along with suggestions for management.

We later asked those involved how they had reached their conclusions. They all immediately recognized SKF as the initials of Smith, Kline, and French, a large drug company. Then, based on their collective experience with SKF products, they recalled as many as they could that are clear, sweet-smelling liquids.

Case 2

A mother called and stated that her child, a 40-pound three-year old, had drunk half a bottle of Tylenol® drops. The expert knew that the active ingredient in Tylenol is acetaminophen and asked the mother to read the label on the bottle to determine how much it had contained. The full bottle had contained one-half fluid ounce, so the ingestion was one-quarter ounce of Tylenol drops. Then the expert consulted a drug handbook to determine the concentration of acetaminophen in Tylenol drops (100 milligrams per milliliter) and used a small hand calculator to compute the total dose (750 milli-

ANTIHISTAMINE/DECONGESTANT POISONING

Antihistamines and decongestants are common ingredients in both prescription and nonprescription products used most often to relieve symptoms associated with allergy or the common cold.

Allergic or immune reactions cause mast cells in the body to release histamine. Histamine receptors in smooth muscle are then stimulated and cause symptoms such as vasodilation, swelling, redness, or running nose. Antihistamines, such as diphenhydramine and chlorpheniramine, block the effects of histamine at some sites. These drugs are described as blocking the H1 histamine receptors. Other agents that block H2 histamine receptors are used clinically to block the stimulant effect of histamine on the secretion of gastric acid. Much of the toxicity that is seen in antihistamine poisoning occurs because of the anticholinergic effects of these agents, that is, their blocking of the neurotransmitter acetylcholine. The blocking produces effects such as dilated pupils, an increase in heart rate, urinary retention, and a drying of secretions, (e.g., dry mouth). Antihistamines also produce effects on the central nervous system. Depression of the central nervous system with drowsiness and possibly coma is common. In some patients, particularly children, central nervous system stimulation occurs. Other central ner-

vous system effects may include hallucinations and convulsions

Decongestants, such as ephedrine, are used alone or in combination with antihistamines to relieve symptoms of the common cold. Their pharmacologic effect is based on their ability to stimulate the adrenergic receptors in nerve endings. The decongestion results from constriction of blood vessels. The most important effects seen in decongestant poisoning are cardiovascular and neurologic. An increase in heart rate and blood pressure is common. The rhythm of the heart may be affected. Central nervous system stimulation with anxiety, nervousness, muscle tremors, and convulsions may occur.

Cases of antihistamine or decongestant poisoning often can be handled easily at home by emptying the stomach through administration of ipecac. Less frequently, hospitalization is required, and more drastic procedures such as gastric lavage or the administration of drugs such as Valium® or physostigmine may be necessary.

In summary, antihistamine and decongestant poisoning can produce potentially life-threatening effects on the cardiovascular and central nervous systems, such as convulsions, hypertension, and arrhythmias. Fortunately, most ingestions are not fatal and are amenable to treatment.

grams). That was divided by the child's body weight (converted to kilograms) to determine the dose per kilogram of body weight (41 milligrams per kilogram). The expert knew that this dose was below the toxic level (100 milligrams per kilogram) and informed the mother that nothing needed to be done; however, the expert did suggest that the child be given something to drink. This suggestion is often made, perhaps to settle the child's stomach, but more likely to settle the mother, who needs to feel that she is doing something useful.

Case 3

A woman called to say that a family member had been stung or bitten by an unknown "bug." When she was asked to describe it, she said that she had never seen anything like it before. It looked like a lady's open-toed slipper; it was "true green" and had a "circle thing like an eye." On both ends it had "things sticking out." The expert immediately said, "That's a saddle-back caterpillar" and assured the caller that it was not dangerous. A moment later another expert entered the room and, amazed that the first had been able to identify the insect so quickly, the person monitoring the calls began to tell him what had happened. She told him of the call and said that the woman said the insect was like nothing she'd ever seen before. Before she could recount the rest of the woman's description, the expert said, "Must have been a saddle-back caterpillar!"

In discussing this interaction with them, we learned that, at the particular season (fall), most reports of weird insects are saddle-back caterpillars. The reason for this is that the caterpillar has a very unusual appearance and most people have never seen one. The first expert said that she immediately suspected a saddle-back caterpillar because the report was of an unusual insect in the fall. She quickly recognized three of the descriptors ("true green," "circle thing like an eye," and "things sticking out") as fitting this hypothesis. The fourth descriptor ("like a lady's open-toed slipper") she discarded because she didn't understand it.

TYPES OF KNOWLEDGE

These three cases, and others, suggest that a number of qualitatively different types of knowledge come into play.

The first case is perhaps the simplest. It is a fact that Smith, Kline, and French manufactures certain drugs, and it is also a fact that only a few of these drugs are clear, sweet-smelling liquids. It is a straightforward task to identify those drugs as candidates for the ingestion.

The second case is only a bit more complicated. Again, the facts are simple: the active ingredient of Tylenol is acetaminophen and the concentration in Tylenol drops is 100 milligrams per milliliter. But certain procedural or algorithmic knowledge is also applied. A simple algorithm is used to compute the drug dose per kilogram of body weight. An algorithmic protocol is used to determine the action to be taken: for doses up to 100 milligrams per kilogram, do nothing; for doses greater than 100 milligrams per kilogram, empty the stomach with ipecac; for doses greater than 200 milli-

grams per kilogram, send the subject to an emergency room and consider administering the antidote Mucomyst®.

The third case is far more complicated. Both formal (saddle-back caterpillars are green) and informal (weird insects in the fall tend to be saddle-back caterpillars) facts are used. But based on these facts, a process of logical inference is followed to reach the conclusion. Both forward and backward chaining seem to be used (see the boxed insert).

Initially, forward chaining occurs: given the set of observations and the season, what insects seem likely? From two initial propositions, "weird" and "fall," and a few rules ("weird" + "fall" imply saddle-back caterpillar), a strong possible identification is deduced. The identification is then taken as a hypothesis to be confirmed by backward chaining. The facts "circle," "green," and "things sticking out" serve to satisfy the goal when appropriate rules are applied ("circle on back" + "green" + "front and rear protuberances" imply saddle-back caterpillar). The hypothesis having been confirmed, with the observation "open-toed slipper" serving neither to confirm nor deny it, that fact is dropped from further consideration.

FORWARD AND BACKWARD CHAINING

Rule-based systems must combine individual rules to reach conclusions. Often, a fact concluded by one rule is used as a premise in another rule. For example, consider the two rules:

- (1) IF it is raining,
THEN there are puddles in the street.
- (2) IF there are puddles in the street,
THEN Johnny should wear his boots.

Suppose we notice that it is raining. Using rule 1, we conclude that there are puddles in the street. We then use rule 2 to infer that Johnny should wear his boots. We have combined, or "chained," rules 1 and 2 to reach this conclusion. When we start from known facts, as we have here, and chain rules to arrive at conclusions, we are using the process called "forward chaining."

In some cases, rather than starting from known facts, we ask questions about possible conclusions. For example, suppose we wish to ask if Johnny should wear his boots. Rule 2 tells us that this will be true if there are puddles in the street. We then ask, are there puddles in the street? From rule 1, we see that this will be true if it is raining. We observe that it is raining, and we therefore answer "yes" to the question "should Johnny wear his boots?" This process, working backward from the desired conclusion, is called "backward chaining."

Some types of logical problem are solved best by forward chaining, some by backward chaining. Many complex problems are best solved by combining both types of chaining.

CHOICE OF ARCHITECTURE

Clearly, a rule-based inference engine is useful in building an expert system. A rule-based language (see the boxed insert) (Texas Instruments Personal Consultant⁷) was selected for the IPECAC system. This language is largely backward chaining, but by using "antecedent rules" a limited amount of forward chaining can be accomplished. The third case analyzed could be implemented easily in the system.

Factual and algorithmic knowledge can also be represented by rules. Indeed, we experimented with this type of knowledge representation early in our work, but found that, while possible, it was inefficient. Other, more natural, ways of representing the knowledge were found. The overall system architecture that was selected con-

sists of the three components, each corresponding to a different type of knowledge, described below.

Factual Knowledge

Factual knowledge is represented and stored in a relational database (see the boxed insert). Thus the relation "analgesic products" might be tabulated as follows:

ANALGESIC PRODUCTS

| | | | | |
|---------------|---------------|-----------|--------|-------|
| Bayer® | aspirin | 325 mg | tablet | white |
| Empirin® | aspirin | 325 mg | tablet | white |
| Tylenol® | acetaminophen | 325 mg | tablet | white |
| Tylenol®drops | acetaminophen | 100 mg/ml | liquid | red |

WHAT IS A RULE-BASED SYSTEM?

Much of what experts know can be formulated into rules. Rules state that IF some condition or set of conditions is true, THEN some logical consequence follows. For example, IF the weather is breezy AND there are no thunderstorms in sight, THEN it is a good day to fly a kite. This rule concludes a fact from its premise. Some rules specify actions to be taken: IF I am flying a kite AND a thunderstorm appears, THEN I should stop flying the kite.

Rules have been incorporated into computer programs for decades, but, in the traditional style of programming, the rules are an intrinsic and inseparable part of the code itself. The programmer must consider carefully how the rules interact; the sequence in which they are used is usually fixed once the program has been written. If the user wishes to change a rule or insert a new one at a later time, he must delve deeply into the structure of the program and usually make changes that are both significant and difficult.

Rule-based systems attempt to simplify the programmer's task by viewing the rules as data to be fed into a program rather than as part of the program itself. A single program of general utility called an "inference engine" is first written that can take a set of facts and rules as input and infer the logical consequences. Thus, the rule "IF someone is a man, THEN that person is mortal" and the fact "Socrates is a man" can be input as data into the inference engine, which will infer that "Socrates is mortal" as its output.

In principle, inference engines allow rules and facts to be entered in any sequence and to be added, deleted, or modified without it being necessary to work through the logical structure of the entire body of knowledge; in practice, this is not true. Currently available inference engines, such as the Texas Instruments Personal Consultant, do much to ease the entry of facts and rules, but the sequence in which rules are encountered and the overall logical structure are still concerns of the programmer.

WHAT IS A RELATIONAL DATABASE?

Relational databases are systems that store information in the form of tables called "relations." Each relation has a name. For example, the relation PEOPLE could look as follows:

PEOPLE

| Name | Sex | Nationality |
|------------|-------|-------------|
| Socrates | man | Greek |
| Napoleon | man | French |
| Helen | woman | Greek |
| Josephine | woman | French |
| Washington | man | American |

Another relation, PROFESSION, could look as follows:

PROFESSION

| Name | Profession |
|------------|-------------|
| Socrates | philosopher |
| Napoleon | general |
| Washington | general |

New relations can be formed from old relations by using logical operations. One such operation is called the "join." If, for example, we wished to have a new relation showing the nationalities of people in each profession, we would join the two relations PEOPLE and PROFESSION by matching the names to get the new relation:

RESULT

| Profession | Nationality |
|-------------|-------------|
| philosopher | Greek |
| general | French |
| general | American |

Many other types of logical operation are possible. Individual pieces of information, such as the fact that Socrates was a Greek philosopher, can be retrieved readily from such a system.

where the first column represents the product name; the second, the active ingredient; the third, the dosage; the fourth, the physical description; and the fifth, the color.

This database could be searched by various standard techniques to find all red liquids or the dose per tablet of a particular drug. There are several hundred commercial products containing antihistamines or decongestants.⁶ At present, only a limited number of products have been entered in the database for demonstration purposes, but there will be no problem in expanding it for the operational system.

Algorithmic Knowledge

Algorithms and protocols are represented by programs written in a procedural language. For example, the acetaminophen protocol can be represented informally by

```
IF
  (DOSE < 100)                "do nothing"
  (DOSE ≥ 100) & (DOSE < 200) "administer ipecac"
  (DOSE ≥ 200)                "send to hospital"
```

In the IPECAC system, algorithmic knowledge is represented by functions written in IQ LISP. The two functions currently implemented are ANTIHISTAMINE-PROTOCOL and DECONGESTANT-PROTOCOL, each of which is considerably more complicated than the example given above.

Inferential Knowledge

Inferential knowledge is represented by IF-THEN rules in the Texas Instruments Personal Consultant⁷ language. Sixty-eight rules compose the pilot system. For example,

```
RULE 026
IF
  (1) antihistamine or decongestant ingestion is suspected,
  (2) an anticholinergic response is present, and
  (3) symptoms include hallucinations,
THEN
  (1) there is strongly suggestive evidence that the ingested drug is known, and
  (2) there is strongly suggestive evidence that the drug name is diphenhydramine.
```

When such rules are encoded into the Personal Consultant language, database queries and calls to procedures in LISP may occur in either the premises or consequences (actions) of rules.

DOMAIN

The present system is capable of handling the initial consultations on cases of antihistamine or decongestant poisoning. Both home and emergency room procedures are supported. Twelve drugs and a number of common products containing the drugs are known to the system. Twelve different symptoms are recognized. The consul-

tation can work from the drug or product name or can attempt to identify the nature of the poisoning from symptoms alone. It can also consider both drug identification and symptoms and determine whether they are consistent. Ten recommendations for intervention can result, ranging from merely observing the patient, through performing a toxicology screen gastric lavage or administering one or more drugs.

APPRENTICESHIP

The system is now undergoing clinical testing at the Maryland Poison Center (see Fig. 2). The clinical staff has been encouraged to use the system, both by experimenting with it and by running actual cases through it, though not in real time.

We expect that rule amplification and fine tuning will be needed; therefore, we regard testing as an inseparable part of the iterative design process. Incorporating criticisms and suggestions of experts into the system will be vital. We call this process the apprenticeship mode of learning.

We are using a questionnaire to record experts' experiences with the system. Each expert is also asked to give an overall subjective rating of satisfaction. Preliminary results have shown that users agree with IPECAC's recommendations about half the time. On a 1-to-5 scale, the average user rating for accuracy has been 2.8, and the average user satisfaction has been 2.6. More importantly, initial user experiences have identified a number of problem areas. For example, at present IPECAC does not use objective criteria for evaluating the seriousness of symptoms (with the exception of tachycardia, for which objective heart rates are used). Rather, the user is asked to state subjectively whether the symptom is more severe than would be expected with a therapeutic dose. Additional rules are needed in this area.

Another problem lies in IPECAC's incorporating the official Poison Center protocols as part of its algorithm-




Figure 2—The IPECAC System is implemented on a Texas Instruments Portable Professional Computer with a 10-mega-byte hard disk.

A DIALOGUE WITH IPECAC

A typical interaction with IPECAC is shown below. IPECAC's questions and recommendations are shown in capital letters, and the user's responses are written in lower case.

IS ANTIHISTAMINE OR DECONGESTANT INGESTION SUSPECTED?

yes

DO YOU KNOW WHAT DRUG WAS INGESTED?

yes

HOW WILL YOU IDENTIFY THE SUBSTANCE INGESTED?

product name

WHAT IS THE NAME OF THE PRODUCT?

benadryl elixir

HOW MANY UNITS (TABLETS, CAPSULES, TEASPOONS OR DROPS) OF DRUG WERE INGESTED?

65

THE QUANTITY OF DRUG INGESTED IS: 1625 MG.

THE NAME OF THE DRUG INGESTED IS: DIPHENHYDRAMINE

ARE SYMPTOMS PRESENT?

yes

WHICH OF THE FOLLOWING SYMPTOMS ARE PRESENT [a list is provided]?

dry mouth, tachycardia, dilated pupils, seizures

WHAT IS THE SUBJECT'S AGE?

child

ARE THE SYMPTOMS MORE INTENSE THAN WOULD BE EXPECTED WITH A THERAPEUTIC DOSE?

yes

WHERE IS THE SUBJECT?

emergency room

WHAT TYPE OF CALL IS THIS?

primary

IS THE SUBJECT DROWSY OR LACKING IN GAG REFLEX?

yes

HAS VALIUM® OR DILANTIN® BEEN ADMINISTERED?

no

RECOMMENDATIONS ARE AS FOLLOWS:

PERFORM GASTRIC LAVAGE ADMINISTER VALIUM® OR DILANTIN®

During the course of this dialogue, the user can ask IPECAC questions about its reasoning. For example, responding to an inquiry from IPECAC with the question "why?" will cause it to print out the rule it is currently trying to satisfy. Sometimes, access to IPECAC's reasoning is even more valuable to the user than are its recommendations.

mic knowledge. The clinical staff does not follow the protocols rigidly. Various criteria, such as the probability of suicide, can override protocol recommendations.

We expect that it will be necessary to add a considerable number of new rules and to fine tune old ones during the apprenticeship phase. On the basis of the criticisms seen so far, we feel there is no serious obstacle to achieving our goal—a system that can accurately handle 75 percent of all inquiries and pass on those cases that are beyond its expertise.

FUTURE GROWTH

There are several ways in which IPECAC can grow, and we intend to pursue them all.

It can grow vertically. That is, we can expand it beyond the initial consultations to more extensive management of complicated cases. This is an obvious direction to pursue, but ultimately we might question its clinical usefulness. Complicated cases are somewhat rare, and they might require more expertise than can be put into our small system. Also, strategically speaking, it is the truly rare and complicated cases that should be reserved for human experts; the simpler, more common cases can be handled by expert systems or paraprofessionals. Nevertheless, we shall try.

More important, perhaps, is horizontal growth. That is, we hope to extend the system to handle initial consultations concerning common substances other than an-

tihistamines and decongestants, for example, other common drugs (analgesics, sedatives), common household substances (bleach, adhesives), plants (poke berries, yew berries), or insects (wasps, bees, but *not* saddle-back caterpillars). Extension in this direction would be truly useful. Such a system, operated by a staff with limited training, could perform initial screening of inquiries, passing on to clinical experts those cases it could not handle.

Finally, we must also consider the user interface. The present menu-driven system can be awkward and slow when it is used in real time. A typical consultation can take a few minutes during which inquiries are made and answered sequentially. Since IPECAC asks many of its questions only when they are needed by the diagnostic process, some sequencing of inquiries is unavoidable, but execution could be made more efficient by using parallel full-screen input of the initial information that is required in all cases. Johannes et al.,⁸ at Johns Hopkins Hospital, have demonstrated the utility of voice input in a clinical setting using an IBM Personal Computer. Voice input is also available for the Texas Instruments

Professional Computer. The combination of voice input and color graphics (especially for pill and plant identification) may significantly increase the power of the system.

REFERENCES and NOTES

- ¹By the American Association of Poison Control Centers.
- ²B. N. Rumack, ed., *Poisindex Information System*, Micromedex Co., Denver (1986).
- ³Maryland Poison Center, internal statistical records.
- ⁴L. M. Haddad and J. F. Winchester, *Clinical Management of Poisoning and Drug Overdose*, W. B. Saunders & Co., Philadelphia (1983).
- ⁵C. H. Cole, ed., *The Harriet Lane Handbook*, Tenth Ed., Year Book Medical Publishers Inc., Chicago (1984).
- ⁶E. K. Kastrup et al., eds., *Facts & Comparisons*, J. B. Lippincott Co., St. Louis (1986).
- ⁷*Personal Consultant Expert System Development Tools User's Guide*, TI Part No. 2243763-0001, Texas Instruments, Inc. (1985).
- ⁸R. S. Johannes, S. N. Kahane, W. J. Ravich, and H. P. Roth, "The Combined Use of Voice and Graphics in a Gastroenterological Data Entry System," in *Proc. American Voice Input-Output Soc.*, pp. 241-249 (Sep 1985).

ACKNOWLEDGMENTS—We gratefully acknowledge the support of Texas Instruments who, through their University Expert Systems Grant program, provided hardware and software for this project. Manpower support was provided by an APL Janney Fellowship.

THE AUTHORS

ANDREW D. GOLDFINGER (left) is a section supervisor of the Computer Science and Technology Group in the Space Department. Born in New York City in 1945, he studied physics at Rensselaer Polytechnic Institute and Cambridge University and received a Ph.D. degree in physics in 1972 from Brandeis University. In 1983, he obtained an M.S. degree in applied behavioral counseling from The Johns Hopkins University. Dr. Goldfinger is currently working on applications of artificial intelligence to remote sensing, image processing, and VLSI design.

GARY M. ODERDA (right, seated) is director of the Maryland Poison Center and an associate professor at the University of Maryland School of Pharmacy. He was born and raised near San Francisco and received a Pharm.D. degree from the University of California at San Francisco in 1972 and a master's degree in public health from The Johns Hopkins University School of Hygiene and Public Health in 1982. His current research interests include clinical toxicology and the epidemiology of poisoning and other injuries. Mr. Oderda has been active in the American Association of Poison Control Centers and now serves as its secretary-treasurer.

RALPH F. WACHTER (right, standing) is a member of the senior professional staff in the Milton S. Eisenhower Research Center at APL and is a scientific officer in the Information Science Division of the Office of Naval Research in Arlington, Va. He graduated



from The Johns Hopkins University with bachelor, master, and doctorate degrees in mathematics. His current areas of interest include software engineering, artificial intelligence, and applied mathematics.